Institute for Artificial Intelligence
Franklin College of Arts and Sciences
UNIVERSITY OF GEORGIA

AI Research Day 2024
AI Transforming

Keynote: Mark Riedl
"The Future of AI is Human-Centered"

Professor, School of Interactive Computing, Georgia Institute of Technology
Associate Director, Georgia Tech Machine Learning Center

Panel Discussion
Lightning Talks
Research Poster Reception

Including awards for best posters

April 22, 2024
2:00 - 7:00 pm
Tate Reception Hall

Contents

Schedule

- 2:00 PM Welcome
- 2:15 PM Keynote
- 3:15 PM Lightning Talks (Session I)
- 3:45 PM Break
- 4:00 PM Anniversary Acknowledgement
- 4:10 PM Panel Discussion
- 5:10 PM Break
- 5:25 PM Lightning Talks (Session II)
- 6:00 PM Poster Session

Organizing Committee

- **Khaled Rasheed** (Organizing Committee Chair; IAI and School of Computing)
- **Soheyla Amirian** (School of Computing)
- **Pete Bettinger** (School of Forestry and Natural Resources)
- **Evette Dunbar** (Event Planning – IAI)
- **Tianming Liu** (School of Computing)
- **Frederick Maier** (IAI)
- **Neal Outland** (Department of Psychology)
- **Ramviyas Parasuraman** (School of Computing)

Keynote: Mark Riedl

@ 2:15 PM

Professor, School of Interactive Computing, College of Computing, Georgia Institute of Technology; Associate Director, Georgia Tech Machine Learning Center.

The Future of AI is Human-Centered

Over the last few years, AI has rapidly moved out of the research lab and into the hands of everyday users. This has been due to technological breakthroughs in deep learning and large language models. However, technologies that work well in the controlled environment of a research lab doesn’t necessarily perform the same out in the real world populated by non-expert users. Human-centered computing refocuses technology on the human by asking: what should our technology look like for it to enhance the human condition, and how so we get there? We will explore some of the ways that AI can become more human-centered from AI explanations to human-gent interaction.

Lightning Talks (Session I)

@ 3:15 PM

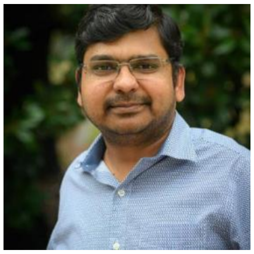

Ramviyas Nattanmai Parasuraman
Assistant Professor, School of Computing

GPS-denied Localization of Mobile Sensors and Robots

I will present my recent work in the research of relative localization in mobile robots and sensors, where the robots collaborate their measurements and apply Bayesian sampling and graph-theoretic approaches to perform relative localization and exploration in a distributed manner. The novelty of these approaches lies in their ubiquity of relying only on minimal sensor data, making them applicable to SWAP-constrained devices like IoT nodes and swarm robots. I will also present some extensions of these works to an autonomous multi-robot exploration domain by presenting a distributed earning technique using ad-hoc map merging and significantly reducing communication and computation costs.

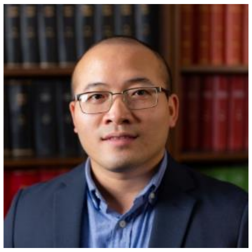

Guoming Li
Assistant Professor, Poultry Science

Transformative AI Techniques for Poultry Production

Poultry production continues to grow to provide affordable animal proteins (e.g., eggs and meat) for humans. It is labor-intensive and time-consuming to inspect the details of individual birds given the challenges of labor shortage, biosecurity, and biosafety in poultry farms. The AI techniques have been increasingly transformed and deployed in the poultry industry to support disease management, welfare assessment, behavior monitoring, housing designs, and robotics. Future research directions of transformative AI in poultry include generalizable AI models, user-friendly assisted AI tools, and biologically-meaningful phenotypic trait extractors.

Juvis Mbeng
Graduate Lab Assistant, Franklin-Physics &Astronomy

Enhancing Clarity of Telescope Images through PCA-based Reconstruction

Interferometry is a telescope technique where signals are combined from multiple receivers in such a way that the traditional telescope resolution limitation is overcome. As such, interferometry is able to resolve either very distant objects, such as the event horizon of a black hole in a distant galaxy, or very small ones, such as planets forming around a nearby star. Telescope images often suffer from noise and other artifacts that hinder the clarity of observations, particularly of forming exoplanets. This project aims to leverage Principal Component Analysis (PCA) for dimensionality reduction and image reconstruction to enhance the clarity of telescope images. By extracting meaningful components and training a model to assign weights to these components, we intend to produce clearer images of planet formation. Doing so is essential to correctly measuring the mass of a forming exoplanet, which is its most fundamental parameter, since it determines if it will be small and rocky, like Earth, or large and gaseous, like Jupiter. Only one of these planets hosts life.

Panel Discussion

@ 4:10 PM

Mark Riedel
Professor, School of Interactive Computing, College of Computing, Georgia Institute of Technology; Associate Director, Georgia Tech Machine Learning Center.

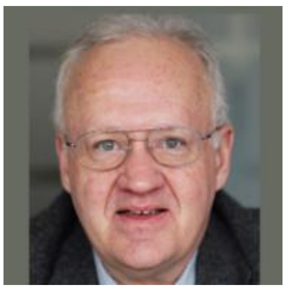

Michael A. Covington
Director of Research for FormFree, Inc.

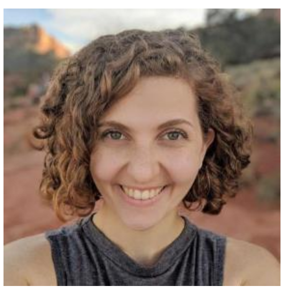

Ari Schlesinger
Assistant Professor, School of Computing

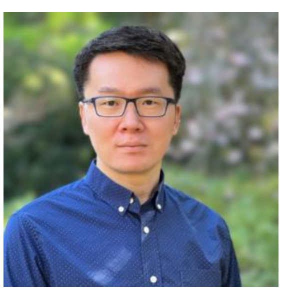

Jin Sun
Assistant Professor, School of Computing

Lightning Talks (Session II)

@ 5:25 PM

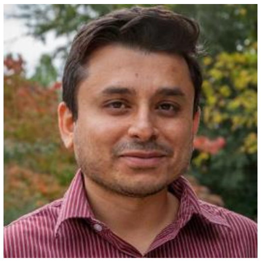

Prashant Doshi
Professor, School of Computing

Agent-based Active Cyber Deception for Adversarial Intent Recognition

Abstract: Deception using honeypots is a tactic of cyber security for confusing the attackers and delaying them by wasting their compute cycles. I will outline a novel approach that uses an intelligent and adaptive agent on a honeypot host. The agent employs deception to engage with a cyber adversary and recognize its intent. We model cyber deception as a sequential decision-making problem in a two-agent context. Sophisticated cyber attacks are primarily orchestrated by human actors. Hence, we also focus on the human factors of the attacker's decision-making process. Humans deviate from rational decision-making due to various cognitive biases. Our approach models and exploits cognitive biases of the attackers in planning for active deception.

Fei Dou
Assistant Professor, School of Computing

Transforming IoT Landscapes: The Role of AI and Machine Learning

Abstract: The integration of Artificial Intelligence (AI) within the Internet of Things (IoT) represents a significant shift towards more dynamic, responsive, and efficient systems across various sectors. This talk will explore the transformative impact of AI on IoT, emphasizing machine learning technologies that enhance the capabilities of IoT systems. We will discuss how AI-driven algorithms enable IoT systems to process and analyze data efficiently, securely, and cost-effectively. Through specific examples, the talk aims to provide insights into the practical applications and benefits of AI in revolutionizing IoT landscapes, paving the way for more innovative and sustainable technological solutions.

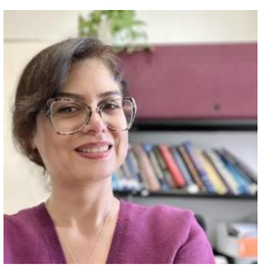

Soheyla Amirian
Lecturer, School of Computing

Transforming Orthopedic Image Analysis using AI: Deep Image Augmentation for Advancing Knee Osteoarthritis Diagnosis

Deep learning strategies have demonstrated very successful applications in medical image analysis, particularly in orthopedics. However, progress in this field has been impeded by the limitation of standardized ground-truth data. To tackle this problem, we aim to introduce a pioneering solution, namely a deep image augmentation pipeline, that copes with this challenge by generating synthetic knee radiographs for training purposes, with a specific focus on Kellgren–Lawrence (KL) grading for knee osteoarthritis. Leveraging innovative techniques, our approach demonstrates the ability to produce high-fidelity images and accurately label them, even with limited training samples. Experimental results demonstrate the efficacy of our proposed computational pipeline, with synthetic images exhibiting remarkable fidelity and achieving high accuracy in labeling. This approach culminates in the creation of a substantial dataset poised for further research and clinical applications. This work highlights the promise of deep image augmentation in augmenting datasets for orthopedic image analysis, offering further research avenues for broader applications in imaging diagnosis.

Poster Session

@ 6:00 PM

Poster #1

Physics-informed machine learning for infectious disease forecasting

Ying Qian, Eric Marty, Avranil Basu, Eamon B. O'Dea, Xianqiao Wang, Spencer Fox, Pejman Rohani, John Drake, He Li

As witnessed in the Coronavirus disease pandemic, accurate forecasting of the spread of contagious illnesses has become increasingly important to public health policymaking and could prevent the loss of millions of lives. To better prepare for the future pandemic, it is essential to continuously improve the capabilities of disease forecasting by introducing novel computational methods. In this work, we propose a new infectious disease forecasting model based on physics-informed neural networks (PINNs), one of the most popular models in the emerging area of scientific machine learning. The proposed PINNs model enables the incorporation of the conventional compartment models into the loss function, thereby assimilating the models and data using neural networks (NNs). This significant feature could prevent the model overfitting, which often occurs when training the machine learning models with observation data alone. In addition, we employ an additional sub-network to account for mobility, vaccination, and other candidate covariates that dictate the transmission rate, the key parameter in the compartment model. To demonstrate the capability of the proposed model, we examine the performance of the model using the state level COVID-19 data in California. Our simulation results show that the model predictions of the number of cases, deaths, and hospitalizations are consistent with the reference. In particular, our results show that the proposed PINNs model outperforms the original compartment model, NN model and the naive model in predicting the dynamics of the pandemic. We also systematically investigate how the accuracy of model predictions is affected by utilizing the length of historical data. Our investigation shows that the model performance changes non-monotonically with the increased amount of the historical data for training the PINNs, implying that redundant historical data could deteriorate the model performance. In summary, our results show that the proposed
infectious disease forecasting model can potentially enhance the capacity to predict the characteristics of COVID-19 pandemics at the state level. Furthermore, the proposed computational framework can be equally applicable to both seasonal epidemics and future pandemics and thus could have a much broader impact.

Poster #2

Transfer Learning on Physics-Informed Neural Networks for Tracking the Hemodynamics in the Evolving False Lumen of Dissected Aorta

Mitchell Daneker, Ying Qian, Eric Myzelev, Arsh Kumbhat, He Li, and Lu Lu

Aortic dissection is a life threatening event responsible for significant morbidity and mortality in individuals ranging in age from children to older adults. A better understanding of the complex hemodynamic environment inside the aorta enables clinicians to assess patient-specific risk of complications and administer timely interventions. In this study, we propose to develop and validate a new computational framework, warm-start physics-informed neural networks (WS-PINNs), to address the limitations of the current approaches in analyzing the hemodynamics inside the false lumen (FL) of Type B aortic dissections (TBAD) vessels reconstructed from apolipoprotein null mice infused with AngII, thereby significantly reducing the amount of required measurement data and eliminating the dependency of predictions on the accuracy and availability of the flow boundary conditions. Specifically, we demonstrate that the WS-PINN models allows us to focus on assessing the 3D flow field inside FL without modeling the true lumen and various branched vessels. Furthermore, we investigate the impact of the spatial and temporal resolutions of MRI data on the prediction accuracy of the PINN model, which can guide the data acquisition to reduce time and financial costs. Finally, we consider the use of transfer learning to provide faster results when looking at similar but new geometries. Our results indicate that the proposed framework can enhance the capacity of hemodynamic analysis in vessels with aortic dissections, with the promise of eventually leading to an improved prognostic ability and understanding of the development of aneurysms.

Poster #3

Automated Data-Driven Discovery of Material Models Based on Symbolic Regression: A Case Study on Human Brain Cortex

Jixin Hou, Xianyan Chen, Taotao Wu, Ellen Kuhl, Xianqiao Wang

We introduce a data-driven framework to automatically identify interpretable and physically meaningful hyperelastic constitutive models from sparse data. Leveraging symbolic regression, an algorithm based on genetic programming, our approach generates elegant hyperelastic models that achieve accurate data fitting through parsimonious mathematic formulae, while strictly adhering to hyperelasticity constraints such as polyconvexity. Our investigation spans three distinct hyperelastic models—invariant-based, principal stretch-based, and normal strain-based—and highlights the versatility of symbolic regression. We validate our new approach using synthetic data from five classic hyperelastic models and experimental data from the human brain to demonstrate algorithmic efficacy. Our results suggest that our symbolic regression robustly discovers accurate models with succinct mathematic expressions in invariant-based, stretch-based, and strain-based scenarios. Strikingly, the strain-based model exhibits superior accuracy, while both stretch- and strain-based models effectively capture the nonlinearity and tension-compression asymmetry inherent to human brain tissue. Polyconvexity examinations affirm the rigor of convexity within the training regime and demonstrate excellent extrapolation capabilities beyond this regime for all three models. However, the stretch-based models raise concerns regarding potential convexity loss under large deformations. Finally, robustness tests on noise-embedded data underscore the reliability of our symbolic regression algorithms. Our study confirms the applicability and accuracy of symbolic regression in the automated discovery of hyperelastic models for the human brain and gives rise to a wide variety of applications in other soft matter systems.

Poster #4

Sub-Cluster Classification of Flow Cytometry Imaging

Subhadeep Sengupta, He Li

Standard clotting tests do not work accurately in studying the mechanisms of thrombosis in COVID-19-affected patients. Flow cytometry offers a more effective way to study sub-cluster cellular components in blood samples. Manual gating is used to analyze and detect sub-cluster components in flow cytometry imaging. Using machine learning, we can automatically cluster cells with the advantage of being less error-prone while saving time to focus on cell cluster interactions and their role in immuno-thrombosis caused by COVID-19. The proposed method splits the process into two phases, with the first phase using a convolutional neural network to robustly classify grayscale channel images into single and multiple cell clusters. The second phase takes the multiple cell clusters and color-classifies sub-cluster components present in the image using existing criteria. The CNN model provides an accuracy of 87.5% on the test set, and the second phase is currently a work in progress.

Poster #5

A Survey of Machine Learning: How Implementing the CIFAR-10 Dataset Prepares Young Computer Scientists for Success

Meg Gleason, Soheyla Amirian

Junior computer scientists can feel overwhelmed with the barrage of information surrounding machine learning, despite a solid mathematical and programming background. Presuming little conceptual background knowledge with adequate technical skills, I explain how a classification task with the CIFAR-10 image dataset is suited for a beginner to delve into the world of supervised learning. An iterative approach is taken with this analysis, beginning with a kNN implementation on raw pixel values and ending with a Deep Neural Network on the reduced dimensionality of the extracted features of the images.

Poster #6

Metamaterial-GPT: A Collaborative Multi-Agen Approach For Directed Metamaterial Design Using Advanced Large Language Models

Jie Tian, Xianqiao Wang

Machine learning-assisted metamaterial design currently remains somewhat isolated. This field typically employs varied strategies for different models and optimization goals, which include not only optimization methods but also model representation techniques. Currently, large language models (LLMs) are transforming various domains; however, the unique aspects of metamaterial design pose significant challenges in adopting these new trends. Despite the strong performance and future potential of LLMs, their application in engineering is limited by issues of stability and controllability, making them less effective than smaller machine learning models. In this study, we aim to improve the usability of LLMs in metamaterial design by applying a multi-agent approach and text-to-image models. We also introduce a framework designed to aid in creating metamaterials of various forms.

Poster #7

Enhancing Patient Safety: A Smartphone-Based Fall Detection Application

Soheyla Amirian, Manoj Tumkur Shivashankar, Lakshmi Meghana Bojja, Rohan Swapneel Intipalli, Ahmad P. Tafti

This work strives to enhance patient safety by developing an Android application utilizing smartphone sensors such as the accelerometer and gyroscope to detect potential falls. The primary objective is to create a user-friendly application capable of identifying probable falls and promptly notifying designated emergency contacts. Upon detecting a potential fall, the app triggers an alert to inform the specified emergency contacts. Key milestones achieved include implementing user login and registration, integrating the app's user interface with a Firebase backend, and storing data using Firebase. Additionally, the project involved detecting smartphone sensor values and calculating potential falls based on this data. Currently, fall detection relies on analyzing sensor data through predefined thresholds and orientation changes, lacking a machine-learning model for more precise detection. The future scope mainly focuses on the clinical validation of the smartphone app and investigating its feasibility and functionalities in real-world clinical settings.

Poster #8

Segmentation of Knee Cartilage in MRIs using Deep Learning: Building a Fully-annotated Imagine Dataset

Nikhita Biradar, A. Patel, M. Dehipitiarachchi, V. Ganjikunta, P. Bodanki, I. A. Siddiqui, L. A. Carlson, M. Abedian, N. Littlefield, A. Kanda, Z. Menezes, A. P. Tafti, J. F., Plate, S. Amirian

This initiative is motivated by the imperative to elevate the clinical impact of knee-related condition diagnosis and monitoring, particularly in cases of osteoarthritis. The overarching goal is to enhance standardization in measurements and assessments across diverse medical facilities, concurrently automating processes to amplify efficiency in MRI analysis and subsequent patient care. The existing challenge lies in the undetected degradation of knee cartilage until advanced stages due to current methodologies being time-intensive, error-prone, and lacking precision in early detection. This research proposes the urgent need for an efficient, accurate, and automated system. Leveraging deep learning networks for knee cartilage analysis emerges as a promising solution, poised to address these challenges and facilitate timely interventions, thereby improving overall patient outcomes.

Poster #9

Advancing Psychosocial Disability and Psychosocial Rehabilitation Research through Large Language Models and Computational Text Mining

Soheyla Amirian, Vedraj Chavan, Ashutosh Kekre, Boby John Loganathan, Punith Kandula, Charles R. Jonassaint, Ahmad P. Tafti

Mental health rehabilitation and psychosocial impediment research continue to be a pressing area in the healthcare domain that needs to be constantly investigated and be analyzed so that the patient's and clinician's outcomes are improved. With the exploration of psychosocial disability area becoming more and more demanding in the academic world the number of publications has been increasing significantly in literature. These journals not only shed light on severe aspects of diagnostics, preventative measures, treatment strategies, and epidemiological aspects, but also provide insights on adverse effects. Therefore, text mining from the huge historical data of Artificial Intelligence (AI) have a great role in appropriately analyzing all these scientific written materials in time, helping the individual scientists of psychological disabilities to understand better and improving how we care for people with these kinds of difficulties. With the use of PubMed library that has all the scientific literature, this study utilizes the high-tech text mining strategies together with word embeddings and large language models (LLMs) to successfully extract invaluable information prompting discovery and advancement in mental health research. It strives to substantially expand the existing database by generation of a massive textual dataset and thereby enable the scientific community to gain insight into current trends in psychosocial rehabilitation and psychosocial disability at large.

Poster #10

Retraining the Phonet Library Using US English

Harsha Veena Tadavarthy, Margaret E. L. Renwick

Our research details adaptation of Phonet Library, a speech technology that utilizes distinctive features to calculate posterior probabilities for phonological classes, by training it on US English. Phonet's posterior probabilities provide a statistical basis for understanding patterned variability in speech, thus bridging the gap between acoustic data and phonological structures. We retrain the Phonet on an English corpus, and we investigate both its precision in phoneme recognition at designated timestamps and the relationship of its distinctive feature probabilities to linguistic expectations for a subset of vowels and consonants. The retrained Phonet library shows accuracy of over 90% in predicting phonological classes. Our analysis demonstrates Phonet's robustness in capturing basic relationships between theoretical natural classes of sounds, and their realization in English, highlighting its utility in the broader context of speech analysis and phonetic research.

Poster #11

Applying large language models to solve partial differential equation problems in physics

Suqiang Ma,He Li, Xianqiao Wang

We propose a novel approach to enhance the problem-solving capabilities of Large Language Models (LLMs) in scientific domains by synthesizing them with Physics-Informed Neural Networks (PINNs). Our hypothesis suggests that integrating LLMs with established scientific machine learning (SciML) models will improve their understanding of mathematical and physical principles, thereby increasing the reliability of their predictions. To validate our hypothesis, we will utilize the DeepXDE library to generate a comprehensive synthetic dataset for training and evaluating these synthesized models. Additionally, our methodology involves applying the X-LoRA technique to train domain-specific LoRA adapters for a variety of physical PDEs, aiming to boost the LLMs' predictive performance in diverse physical sciences. This initiative paves the way for creating a computational pipeline that fuses SciML models with LLMs to unlock efficient and accurate solutions for scientific challenges.

Poster #12

ImitationBT: Imitation Learning for Behavior Tree Generation from DRL Agents

Shailendra Sekhar Bathula and Ramviyas Parasuraman

Behavior Trees (BT) stand as a favored control architecture among game designers and robotics experts, prized for their modularity, reactivity, and hierarchical structure. These properties enable BTs to offer scalable and clear-cut solutions to a wide array of decision-making challenges. In contrast, Deep Reinforcement Learning (DRL) has demonstrated exceptional performance but faces hesitancy in high-stakes domains due to its reliance on neural networks, which present challenges in verifiability. In this context, we introduce a novel framework designed to bridge the gap between the high performance of DRL and the desirable transparency and verifiability of BTs. By employing imitation learning to capture and transfer the expertise of a reinforcement learning model, we pave the way for generating BTs that are not only effective but also transparent, interpretable, and readily verifiable for real-world problems.

Poster #13

PhysicsAssistant: An LLM-Powered Interactive Learning Robot for Physics Lab Investigations

Ehsan Latif, Ramviyas Parasuraman, and Xiaoming Zhai

Robot systems in education can leverage Large language models' (LLMs) natural language understanding capabilities to provide assistance and facilitate learning. This paper proposes a multimodal interactive robot (PhysicsAssistant) built on YOLOv8 object detection, cameras, speech recognition, and chatbot using LLM to assist students' physics labs. We conduct a user study on ten 8th-grade students to empirically evaluate the performance of PhysicsAssistant with a human expert. The Expert rates the assistants' responses to student queries on a 0-4 scale based on Bloom's taxonomy to provide educational support. We have compared the performance of PhysicsAssistant (YOLOv8+GPT-3.5-turbo) with GPT-4 and found that the human expert rating of both systems for factual understanding is the same. However, the rating of GPT-4 for conceptual and procedural knowledge (3 and 3.2 vs 2.2 and 2.6, respectively) is significantly higher than PhysicsAssistant (p $<$ 0.05). However, the response time of GPT-4 is significantly higher than PhysicsAssistant (3.54 vs 1.64 sec, p $<$ 0.05). Hence, despite the relatively lower response quality of PhysicsAssistant than GPT-4, it has shown potential for being used as a real-time lab assistant to provide timely responses and can offload teachers' labor to assist with repetitive tasks. To the best of our knowledge, this is the first attempt to build such an interactive multimodal robotic assistant for K-12 science (physics) education.

Poster #14

Zero-Shot Image Segmentation for Monitoring Thermal Conditions of Individual Cage-free Laying Hens

Mahtab Saeidifar, Guoming Li, Lilong Chai, Ramesh Bist, Khaled M Rasheed, Jin Lu, Ahmad Banakar, Tianming Liu, Xiao Yang

Body temperature is a critical indicator of health and productivity in egg-laying chickens. The gold standard method for measuring the cloacal temperature of chickens involves using thermometers. However, this method is time-consuming, labor-intensive, and causes stress to birds.

The objective of this research was to apply zero-shot computer vision techniques to automatically segment chickens from thermal images without human intervention. Subsequently, we aim to extract body temperature statistics, such as mean and median, among others, to evaluate the thermal profile of a chicken's body more comprehensively.

Poster #15

Incremental Left-Corner Generative Dependency Parsing as a Model of Language Processing

Donald G. Dunagan, Dustin A. Chacón, John T. Hale

In 2024, large Transformer language models (LLMs) play a central role not only in artificial intelligence, in general, and natural language processing, in particular, but also in computational, psycho-, and neurolinguistics. Recent work, however, demonstrates that LLMs cannot explain human sentence processing difficulty. We aim to better model human sentence processing difficulty by proposing and benchmarking a language model with cognitively realistic inductive biases. Specifically, we propose an incremental left-corner generative dependency parser. Diverging from previous transition-based methods in dependency parsing, which were inherently bottom-up, our system will be mildly predictive (like the human sentence processor) via a left-corner transition system. The parser will also be generative. That is, it will be able to be used as a language model while still explicitly representing grammatical structure. The developed model will be benchmarked against observed human reading time data, will be benchmarked as a dependency parser on the Penn Treebank, and will be used to investigate observed electroencephalography brain data collected as participants read short sentences.

Poster #16

Human-Robot Safety, Social Compatibility, and Interaction Transparency

Tyson Jordan, Prashant Doshi

UGA's THINC Lab is currently recruiting students to participate in our study, "Human-Robot Safety, Social Compatibility, and Interaction Transparency." In our experiment, pairs of subjects participate in a simulation of a medical evacuation, each taking on the role of "casualty" or "bystander". We employ surveys to evaluate our subjects' perceptions of safety while interacting with the robot. Furthermore, we introduce the Human-Machine Teammate Inventory (HMTI), a rigorously validated scale designed to measure human perception of a machine teammate.

Poster #17

Exploration of the Genesis of Language via Artificial Life Simulation

Austin Downes, Morgan Riley, Khaled Rasheed

The genesis of language is a critical topic across numerous disciplines. The aim of this project is to provide insight into hypotheses on this topic by simulating agent populations that, with no initial linguistic knowledge, develop language or language-like systems under evolutionary pressures. Our simulation design is twofold, consisting of environment and agent models. The agent model can be understood as the combination of four logical components: an environment interface and language production and reception, semantic memory, and decision making modules.

Poster #18

Enhancing Clarity of Telescope Images through PCA-based Reconstruction

Juvis B. Mbeng, Cassandra Hall, Lia Medeiro

Interferometry is a telescope technique where signals are combined from multiple receivers in such a way that the traditional telescope resolution limitation is overcome. As such, interferometry is able to resolve either very distant objects, such as the event horizon of a black hole in a distant galaxy, or very small ones, such as planets forming around a nearby star.

Telescope images often suffer from noise and other artifacts that hinder the clarity of observations, particularly of forming exoplanets. This project aims to leverage Principal Component Analysis (PCA) for dimensionality reduction and image reconstruction to enhance the clarity of telescope images. By extracting meaningful components and training a model to assign weights to these components, we intend to produce clearer images of planet formation. Doing so is essential to correctly measuring the mass of a forming exoplanet, which is its most fundamental parameter, since it determines if it will be small and rocky, like Earth, or large and gaseous, like Jupiter. Only one of these planets hosts life.

Poster #19

Integration of Multi-Robot Systems in Heterogenous Swarm Robotics: Building a Robust and Adaptive Testbed

Neil Patel Samuel Lemus, Ramviyas Parasuraman

This research focuses on the integration of multiple robotic systems within a heterogeneous swarm robotics framework, highlighting the development of a versatile testbed. Specifically, integrating the existing SwarmV3 robots with the newly developed HeRo Cube robots with unique docking and charging capabilities. The core achievement of this research is the design and realization of a robust, adaptive testbed that supports the dynamics of heterogeneous swarm robotics. The testbeds versatility allows for extensive experiments such as studying swarm behaviors, robot interactions, and collective task allocation. The successful of multi-robot systems utilizes an overhead camera system paired with AprilTag detection. This project achieves inter-robotic communication strategies that is pivotal for conducting accurate swarm experiments, enabling real-time adjustments. The successful integration of the HeRo Cubes is a significant advancement in swarm robotics by allowing docking capabilities. This enables future research in energy distribution where docked HeRo Cubes could share power resources, enhancing the operational efficiency of the swarm platform. Future swarm experiments can include Task Allocation and Role Differentiation based on energy levels, with HeRo Cube robot acting as mobile charging stations. Other swarm experiments can include Adaptive Path Planning with Energy Constraints.

Poster #20

An Adaptive Method Stabilizing Activations for Enhanced Generalization

Hyunseok Seung, Jaewoo Lee

We introduce AdaAct, a novel optimization algorithm that adjusts learning rates according to activation variance. Our method enhances the stability of neuron outputs by incorporating neuron-wise adaptivity during the training process, which subsequently leads to better generalization---a complementary approach to conventional activation regularization methods. Experimental results demonstrate AdaAct's competitive performance across standard image classification benchmarks. We evaluate AdaAct on CIFAR and ImageNet, comparing it with other SGD and Adam. Importantly, AdaAct effectively bridges the gap between the convergence speed of Adam and the strong generalization capabilities of SGD.

Poster #21

Interactive Team Environment – Robot Trust in Humans

Anusha Challa

When humans collaborate, they form positive or negative experiences with each other. These experiences depend on various factors such as the individual's skills, abilities, and agency. In this paper, we consider human-robot collaborations and present a novel model of an autonomous robot's trust in humans based on the probability of the robot having a positive experience with the human. The model defines a dynamic trust-building process that translates into a computationally accessible implementation. We hypothesize predictors of a positive experience with human teammates and derive trust in individual humans. As the interactions increase, team members develop affinity toward each other. The robot's affinity towards humans can be viewed as kinship, and we also investigate how kinship affects trust and distrust. We present an algorithm for how the robot may use the kinship-mediated trust in its decision-making and demonstrate its use in simulated missions truly requiring human-robot collaboration.

Poster #22

Modeling Human Activity Recognition from Ambient Sensors in Smart-Home Environments

Reshma Seby John

The rise of Ambient Assisted Living (AAL) and machine learning with ambient sensor data highlights the need for accurate human behavior modelling. Existing end-to-end machine learning excels at extracting high-level features from raw signals, but there is a lack of studies exploring the proper encoding of time-series firings from binary environment sensors. These event-triggered sensors with sporadic recordings pose a challenge for in-home activity recognition. This project aims to bridge this gap by developing robust machine learning models for activity recognition using such sensor data.

Poster #23

Automatic Forest Stand Delineation with AI, using Optical Imagery and LiDAR Data

Advait Sankhe, Can Vatandaslar, Frederick W. Maier, Khaled Rasheed, Jiayi Tu

A forest stand is a contiguous community of trees that is often uniform in species composition and biophysical structure to distinguish it from adjacent communities. Management decisions are generally applied to stands of trees rather than individual trees in the field of forestry. Therefore, forest management necessitates accurate and up-to-date stand boundary maps, typically represented as vector polygons, to inform decision-making processes. The traditional manual creation of these maps is based on visual assessments of aerial imagery and is time-consuming, especially for large and complex forested landscapes. This study employs artificial intelligence (AI) techniques and diverse data sources to automate the forest stand mapping process within the Talladega National Forest, Alabama, which covers 93,694 hectares.

Poster #24

Glydentify, a deep learning tool for classifying glycosyltransferase function

Aarya Venkat, Ruili Fang, Zhongliang Zhou, Shaan Gill, Natarajan Kannan

Protein language models have emerged as a powerful tool for predicting protein function by capturing the underlying grammar and syntax of protein sequences. Here, we introduce Glydentify, an open-source and user-friendly application that uses protein language models for the classification of glycosyltransferases (GTs) and donor prediction. Utilizing the state-of-the-art ESM2 protein language model, Glydentify extracts high-dimensional sequence embeddings to accurately classify GTs into fold A families with 92\% accuracy. The tool also predicts GT-A donor binding preferences with an accuracy of 91\%. Notably, Glydentify identifies key residues that contribute to a prediction, thereby adding an explainable component to the application. With an intuitive interface powered by Gradio, Glydentify requires no programming experience from the user, democratizing access to cutting-edge deep learning technologies for GT research. The application is freely available on GitHub and can be accessed directly through any web browser (https://huggingface.co/spaces/arikat/Glydentify).

Poster #25

Incorporation of AI into Cultural Heritage 3D Modeling Workflows

Ella Ma, Mitchell Riggins, Chandler Phillips, Sergio Bernardes, Allison Howard

Preservation of cultural heritage has long been a critical endeavor, aiming to safeguard artifacts and sites of historical significance for posterity. One method of preservation involves the creation of digital twins, integrating technologies such as drone photogrammetry and 3D modeling. However, this is still a time and labor-intensive process, and as advancements in technology - such as with 3D generative AI - develop, new workflow practices are continuously being assessed. Using the Cobb's World project as a case study, which focused on the demand for innovative approaches using technology in heritage preservation, we analyze the feasibility of integrating various forms of AI into the existing methodology.

About the AI Institute

In 1984, interested researchers at the University of Georgia formed the *Artificial Intelligence Research Group*. The Group was officially established as the *Artificial Intelligence Center* in 1994, and it became the *Institute for Artificial Intelligence* in 2008. It is part of the Franklin College of Arts and Sciences.

The mission of the Institute to encourage and support interdisciplinary research in Artificial Intelligence involving University faculty and staff, and to facilitate the integration of Artificial Intelligence technology by industry and government agencies in the State of Georgia and the Nation.

The Institute is both a research and an instructional unit. It has for many years administered both the AB in Cognitive Science and the MS in Artificial Intelligence degrees, and in 2022 a PhD in Artificial Intelligence was approved by the Board of Regents of the University System of Georgia.

Participation in the AI Institute by UGA faculty is voluntary. There are currently 70 associated AI Faculty Fellows, taken from many departments across multiple colleges and schools at UGA. Some of the units involved include: The School of Computing, Philosophy, Linguistics, Psychology, Public Administration & Policy, Political Science, Educational Psychology, Statistics, Theatre & Film Studies, Marine Sciences, Electrical and Computer Engineering, Business Administration, Advertising & Public Relations, English, Geography, Management Information Systems, Epidemiology & Biostatistics, Mathematics and Science Education, Health Promotion & Behavior, Epidemiology & Biostatistics.

More information on the AI Institute can be found at *https://ai.uga.edu*.